arXiv:1610.05819v1 [cs.IR] 18 Oct 2016

Finding Representative Points in Multivariate Data Using PCA

Ashwinkumar Ganesan
University Of Maryland Baltimore County (UMBC)
Baltimore, MD, USA
mandarh1@umbc.edu

Tim Oates
University Of Maryland Baltimore County (UMBC)
Baltimore, MD, USA
gashwin1@umbc.edu

Matt Schmill
University Of Maryland Baltimore County (UMBC)
Baltimore, MD, USA
oates@cs.umbc.edu

Abstract

The idea of *representation* has been used in various fields of study from data analysis to political science. In this paper, we define *representativeness* and describe a method to isolate data points that can represent the entire data set. Also, we show how the minimum set of *representative* data points can be generated. We use data from *GLOBE* (a project to study the effects on Land Change based on a set of parameters that include temperature, forest cover, human population, atmospheric parameters and many other variables) to test & validate the algorithm. *Principal Component Analysis* (PCA) is used to reduce the dimensions of the multivariate data set, so that the *representative* points can be generated efficiently and its *Representativeness* has been compared against *Random Sampling* of points from the data set.

1 INTRODUCTION

The study of environmental sciences and especially *climate change* has grown in importance over the last two decades. Some reasons to study these areas are to understand their effects on human populations and on natural resources available on the planet. They encompass a host of issues that are related to diverse areas, including ecology, government policies, public health and economic policies Ellis [2011]. Global climate change study attempts to understand (and predict) how the climate evolves based on current and historical evidence. It explores the change in climate occurring due to natural phenomenon and anthropogenic reasons. Our effort is to find a way to globally utilize these studies by measuring similarity.

1.1 Land Change Science & Globe

Land Change Science (LCS) is a part of global climate change research and studies the effect of human activity on land and on the climate. As part of LCS, studies of specific geographic locations are conducted across the entire world. The locations where these studies are conducted are based on

a number of factors that are decided depending upon the goals of the study that could be social or ecological. Goals that are social could include understanding the nature of human interaction with the environment and the impact of the change on human lives. Ecological studies try to measure parameters such as deforestation and its effect on the environment as a whole. The size or scale of a study could be based on the amount of funds available or to perform a study for measurement of global climate change (like an IPCC study) that can be used for making policy on Climate Change [2000]. These studies are conducted by scientists on-site or by remotely monitoring these locations through stations that are constructed on the site.

Globe is a global correlation engine to study land change science Ellis [2011]. It tries to fulfil the basic requirement of providing the ability to generalize LCS studies. It has multiple objectives which include constructing of an online social media system to facilitate collaborative work between land change scientists, providing a set of tools to analyse and generalize local observations to larger regions or the entire globe, and constructing scientific workflows Ellis [2011].

1.2 Motivation

Land change studies are expensive. Hence the main motivation is to provide an understanding of how the conclusions of these studies can be generalized to other parts of the world. The notion of how generalized the studies are can help land change scientists to:

1. Reduce the number and cost of studies conducted in the future.
2. Select new locations in the world where studies have not been conducted.
3. Analyse existing studies in a different manner by showing a global pattern of the distribution of a set of selected parameters.

1.3 Representativeness

There are a set of parameters or variables whose information is available for all parts of the world. They are the set of *global* variables. These variables include potential vegetation, land utilized for agriculture, temperature, and many others. Scientists studying land change select a case study or a set of case studies which have been conducted. They select the variables which represent the environment in which the studies were conducted or the variables for which they want to analyze the case study. To generalize the results of these studies, they use these parameters to find other parts of the world that are *similar*. The *representativeness* of a given set of case studies is defined as the extent to which they *cover* the rest of the world. For example, consider a scientist analyzing the Van Vliet Study van Vliet et al. [2012] on trends in swidden cultivation. She would perform the following steps

1. Create a collection of all the locations that are part of study.
2. Select this *collection* from *Globe's* User Interface.
3. Select parameters such as %crops, %tree cover, temperature, Market Access Index and Potential Vegetation as the meta-study uses these variables.
4. Compute representativeness for the collection. Display it on a world map to show which other regions of the world are similar or to what extent the rest of regions are dissimilar to the ones in the study.

1.4 Challenges & Contributions

The number of distinct geographical regions across the world into which the world map can be divided is very large. Each region is has its own global variable information. There are a large number of dimensions which, coupled with the number of sites, makes the representativeness computation expensive. The users (scientists in this case) select a subset of variables in realtime. The results for this calculation are to be provided in realtime. There are essentially two challenges here:

1. To find a method to measure representativeness.
2. To provide an algorithm, to find a minimum set of new locations where studies can be conducted so as to maximize representativeness for the selected set of variables.

Our contribution in this thesis is to provide a mathematical formulation for representativeness, reduction of computation time and increase in efficiency by using a dimension reduction technique and providing a method to validate the results of our selected locations. Dimension reduction techniques are required since the number of dimensions is large and methods that use the original set of dimensions, such as clustering, will not be able to perform computation in realtime. The dimension reduction method we apply is Principal Component Analysis (PCA). PCA is useful because it prioritizes dimensions with higher variance. We show in this thesis how the algorithm can be used to select new locations and the correlation between the geographic points in the original space and single dimension PCA space. This correlation helps us measure the effectiveness of PCA and selected locations. Representativeness is displayed in the form of a *Heat Map* using Google Maps. The map has markers that show the location where each case study was conducted.

2 REPRESENTATIVENESS

2.1 Definition

Representativeness describes how a data point or a set of data points can be used to generalize to the rest of the data set. In case of the Globe project, a data point refers to a specific region in the world where the study is conducted. Consider a distribution of data points where

- D is the given dataset of points
- S is a sample set of points such that $S \subseteq D$
- H is a histogram based on D
- *Bin(H,s)* is the bin where the data value s falls in H
- *P(H,i)* is the height / probability of bin i in histogram H
- All unique bins are defined in a set $B = \{b|\forall_{s \in S} b = bin(H,s)\}$

We define *representativeness* R of a sample set S for a given global dataset D as

$$R(S|D) = \sum_{b \in B} p(H, b) \tag{1}$$

where $0 \leq R(S \mid D) \leq 1$. When a sample set has higher representativeness, then R reaches 1.

The definition of representativeness R is theorized for a dataset D which has a single variable or attribute for each data point. The histogram for the dataset D gives us the frequency of data points in each bin that is defined. Once we know which data points fall in which bin and where points in sample set S lie in the histogram, we know which data points are represented by S. These data points are in set B that is the set of bins b where sample points in S lie (i.e., $b \in B$). The data point $x \in D$ is said to be represented by a sample point $s \in S$, when a certain pre-defined criteria is fulfilled.

Thus representativeness can be explained as a fraction of the total number of data points that fall within a predefined threshold criteria for atleast one of the points in the sample set. All data points within the threshold are completely represented by one of the sample points in S. If a data point falls within the threshold criteria of multiple sample set points, then it is represented by the sample set point where criteria is optimal. In Globe, representativeness shows the fraction of the total land surface on earth that are similar to the locations that are part of a case study and have been studied based on a specific set of parameters.

If we use a multivariate dataset with m dimensions, then the criteria used are modified to consider m dimensions. For example, if the criteria is based on Euclidean Distance, then distance in a single dimension would be

$$d = |(s - x)| \tag{2}$$

where

- d is the distance
- s is a sample point such that $s \in S$
- x is a data point such that x $\in$ D

The distance formula for *m*-dimensional data points would be

$$d = \sqrt{\sum_{i \in m} (s_i - x_i)^2} \tag{3}$$

The distance between the sample point $s \in S$ and *x* shows how close the data point is to the sample point. As $d \to 0$, the data point is considered to be closer. Representation of a data point by a sample point is inversely proportional to the distance. Hence, representation *r* is defined as

$$r(x|s) = |1 - d_x| \tag{4}$$

r is thus a value between 0 and 1 (and maybe greater than 1 in some outlier cases). A scale is created from 0 to 1 and the data points are assigned to each section of the scale (forming histogram *H*). Representativeness *R* is taken as the proportion of the total number of data points that are there in the first scale between 0 and 1. This is because the representativeness provided in the definition is for a binary scale where as represetativeness can be in degrees (like in the case of a heat map explained in the next chapter).

The histogram *H* can be of 2 types: equal probability and equal area as shown in the diagram below where *x* is a single variable and *p(x)* is the probability distribution function (pdf).

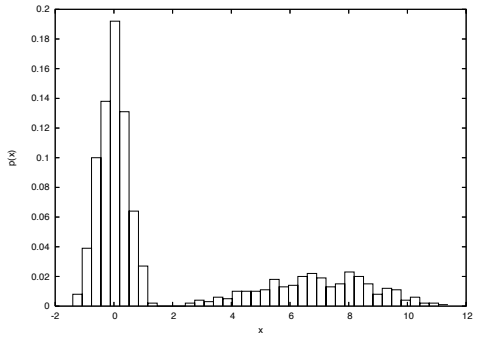

Figure 1: *A histogram of a sample of data from a distribution in which bins have equal area.*

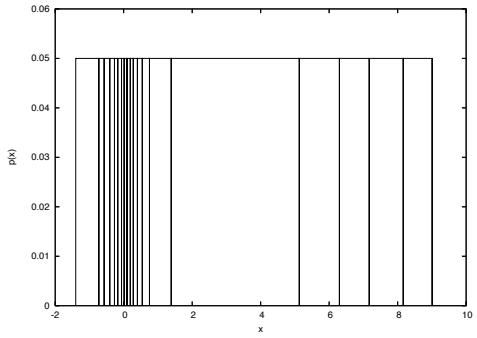

Figure 2: *A histogram of a sample of data from a distribution in which bins have equal empirical probability.*

An equal area histogram is one where the data dimension is divided into bins of equal size (Figure 1). Thus each bin contains the points that fall within a certain range in the given dimension. Representativeness *R* in such a histogram can be maximized by choosing mode points. An equal probability histogram is where the number of points in each bin is equal. Hence the size or width of each bin changes according to the density of the points (Figure 2). Since the bins are equiprobable, points can be selected from any random bin to represent the entire points in that bin.

2.2 Kernels

A kernel function is a function that maps a point onto a scale and is denoted by *K(s - d)*. In the case of a histogram, the kernel function implemented is a step function. It is a set of bins whose points match the threshold criteria (like maximum distance) such that any point in any of the bins represents the all the other points. When the histogram is equiprobable, the step function is used to maximize the number of bins that are part of within the function's limits. The step function is shown in the diagram below.

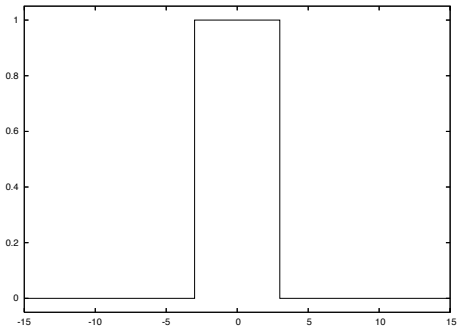

Figure 3: *Sample Step kernel function.*

3 RELATED WORK

In this section, we discuss various methods that can be used to maximize represetativeness. The main aim of the methods described below is to find an optimal set of sites or points such that representativeness can be maximized.

3.1 Clustering

Clustering techniques are a set of methods to group data points that are similar together Jain et al. [1999]. These characteristics of the groups are defined by a pattern of values in their variables. Clustering is an unsupervised learning method. It does not require a training data set to create a model. The groups in which the data points are to be classified need not be known at the start. Hence, clustering can be used for exploratory data analysis to identify patterns in the data. Clustering is a three stage process

1. Extract features from the given set of points.
2. Perform similarity measurement between data points.
3. Create groups based on the similarity measurement.

Clustering techniques are of different types and mainly divided into 2 categories:-
Hierarchial Clustering Techniques - These techniques create groups of points which are similar to each other. Once a group is formed, it creates the next level by combining groups that are similar. In this way, a hierarchy of groups is created with all groups merged at the top most level of the hierarchy. The structure is called a dendogram Jain et al. [1999].
Partitional Clustering Techniques - These techniques try to create a single partition in the dataset as compared to a dendogram which may have a high computation time. The problem occurs when the size of the dataset is large. Partitional techniques try to optimize a certain function based on which the partition is made. Calculating the optimal set of values for the function could again be computationally expensive. Hence an approximation is calculated by executing the algorithm multiple times on the same dataset until the function reaches a state that is *close* to optimal. For example, using squared error a as function to create partitions Jain et al. [1999]. The algorithm is executed until the squared error is reduced to a value that is below a certain pre-determined threshold.

3.2 K-Means Clustering

The k-means clustering algorithm is a widely used algorithm Vattani [2009]. This is a centroid or partition based clustering technique. The algorithm clusters all the data points into *k* clusters. The algorithm starts by selecting an arbitrary set of centroids $c_1, c_2...c_k$. It then assigns each point to the closest centroid c_i. Once the points are clustered, it calculates the center of mass for each cluster to get a new set of centroids. The previous steps are then repeated for the new centroids. After each iteration the set of centroids moves closer to the final set such that the next iteration does not change the set of centroids chosen. This means the center of mass for the *k* clusters calculated remains constant. The algorithm stops computing after this point. The worst case time complexity is $O(n^{kd})$ Vattani [2009] where *n* is the number of data points, *k* is the number of clusters and the points are in a *d*-dimensional space.

3.3 Nearest Neighbor Clustering

This is a hierarchial clustering technique. In this clustering method, the nearest neighbor to each data point is found and the point is assigned to that cluster. A Voronoi decomposition of the data points is performed Bubeck et al. [2009]. There is a threshold or quality function Q_n to put a threshold on the distance that is considered between the point and the cluster. Thus all the points are put into *k* clusters where *k* is user-defined. The clustering is implemented using a graph based structure. Whenever a point closest to the current point is found, an edge is created between them thus linking them in the same cluster Jain et al. [1999]. It is also called agglomerative single-link clustering technique and has a time complexity of $O(n^2)$ Manning et al. [2008].

K-means clustering and nearest neighbor clustering can be used to find a set of *k* centroids that maximize representativeness. K-means clustering generates *k* clusters with unique centroids that are the representative points. For nearest neighbor clustering, we can select any point randomly from each of the *k* clusters generated (as all clusters adhere to the quality function Q_n), as they represent the other points within the cluster.

3.4 Dimension Reduction Techniques

Consider a data set where each point has a large number of variables. These variables may have different scales of values, and different densities and variances. There are a number of possible problems with high dimensional data Carreira-Perpinan [1997a]:

1. Processing high dimensional data (especially when the number of data points is large) is expensive.
2. Even though the number of dimensions is high, the data could be classified or clustered using a smaller subset of variables.
3. As the number of dimensions increases, the values for some variables may become sparse. This is known as the *empty space problem* Carreira-Perpinan [1997a].
4. The *Curse of Dimensionality* states that the number of sample points required to approximate a function increases exponentially as the number of variables / dimensions increases.

A dimension reduction technique is a transformation which reduces number of dimensions required to represent a sample. The reduced set of dimensions may be a subset of the original set of dimensions (for example, using information gain) or could be a completely new set of dimensions. Some of standard dimension reduction techniques that can be used to transform a high dimensional data set are Principal Component Analysis (PCA), and Self Organizing Maps (SOM)Carreira-Perpinan [1997a]. Neural Networks with GIS have also been used for constructing a *Land Transformation Model* which tries to forecast how usage changes Pijanowski et al. [2002]. Self Organizing Maps have been used to perform environmental assessment of regions, grouping based on environmental conditions, and finding out which areas might deteriorate in the future Tran et al. [2003]. Once a SOM is trained, the *k* nodes from the weight vector can be used as centroids representing their respective clusters. As the nodes may not be actual data points, the point closest to each node will be used as a representative point. Training a SOM may require updating the weight vector over several iterations of the data set. The time complexity of a SOM is $O(|E| \bullet (|E| + |V|))$Brugger et al. [2008].

3.5 Clustering Using a Combination Of Methods

Hoffman et. al.Forrest M. Hoffman and Oglesby [2008] use *Multivariate Spatio-Temporal Clustering (MSTC)* to calculate representativeness of sampling networks. MSTC can be performed using a combination of PCA and K-Means clusteringForrest M. Hoffman and Oglesby [2008]. The data set considered is high dimensional and is assumed to contain a lot of redundant information. Hence the method involves reducing the number of dimensions using PCA at the beginning and then performing standard k-means clustering. Hoffman et. al. also provide a set of improvements for performing PCA and k-means clustering. The time required to perform k-means clustering is reduced by decreasing the number of distance computations between the centroid and the other points, based on cluster created and new distances computed. The time complexity of PCA computation is reduced by parallelizing it. The summation of all euclidean distances from points to their nearest sample locations or centroids, is used to measure representativeness of the sample set. Higher the sum, lower is the representativeness of the sample set.

4 CALCULATING REPRESENTATIVENESS

4.1 Principal Component Analysis (PCA)

Principal Component Analysis (PCA) is a technique used to analyze multivariate datasets and to find patterns in the data. It is used as a dimension reduction technique where higher dimension data points can be projected onto a lower dimension space Abdi and Williams [2010]. It takes a data set of n dimensions and projects it onto a set of new dimensions (which again can be a maximum of n) that are orthogonal to each other (although all n need not be used). The new dimensions are not correlated with each other. They are called *Principal Components*. The principal components are calculated by performing an eigen decomposition of the covariance matrix. The covariance between dimensions in the original dataset is used because those dimensions which have a higher variance, provide more information regarding the nature of the dataset Carreira-Perpinan [1997b]. Another reason to use PCA is because it is easy to understand and efficient algorithms exist to compute each step Carreira-Perpinan [1997b]. It is also a widely used method because it is able to reduce noise and the dimension reduction performs data clustering Ding and He [2004]. As we shall see, the cost of computing distance is also reduced as the final distance is in a single dimension.

4.2 Measuring Representativeness Using PCA & Histograms

The following section provides the details of the algorithm to calculate Representativeness of a given sample set of regions and how to select a new set of sites to improve representativeness across the entire world for the selected set of attributes or variables. Land Change scientists know many regions across the world are not represented by sample sites in their case studies by design. In *Globe*, they are provided the option to filter out these regions, so that they are not considered as part of the analysis. Hence we apply our method to the list of unfiltered regions.

The following diagram shows the algorithm's workflow:

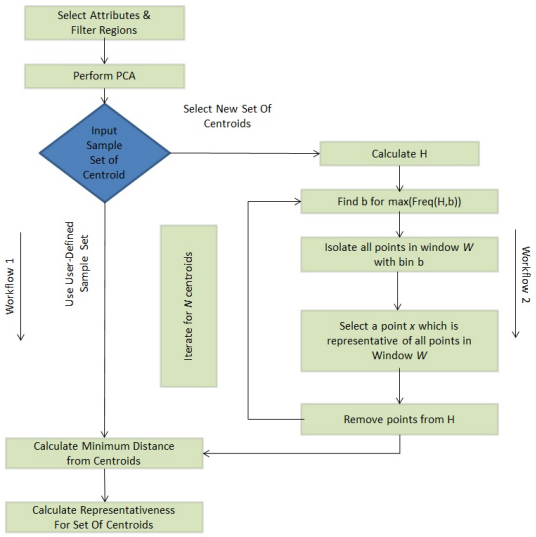
Select Attributes & Filter Regions
Perform PCA
Input Sample Set of Centroid
Select New Set Of Centroids
Calculate H
Find b for max(Freq(H,b))
Isolate all points in window W with bin b
Select a point x which is representative of all points in Window W
Remove points from H
Workflow 1
Use User-Defined Sample Set
Iterate for N centroids
Workflow 2
Calculate Minimum Distance from Centroids
Calculate Representativeness For Set Of Centroids

Figure 4: *Workflow to Calculate Representativeness*

As shown in the diagram above, the algorithm has two distinct workflows to solve two objectives:

1. The first workflow calculates the representativeness for the user-defined Sample Set of regions where case studies were conducted.
2. The second workflow provides a method to select *N* regions to maximize representativeness for the user-defined set of attributes. These are the places where the system can recommend that a study could be conducted. Once the regions are selected the representativeness of these sites is calculated.

Both workflows use a distance metric to calculate how close a location is to another location where a study was conducted.

$$FD(x) = min(D_p(x) - D_p(s)_{s \in S}) \tag{5}$$

where x is a specific location across the world and s is a place in the world where a study was conducted. s is part of a larger Sample Set S. D_p is the projection of a location (or data point) onto the first principal component. The final distance $FD(x)$ is the minimum distance between x and S.

Representativeness of the Sample Set, as given in equation 1 can now be transformed to the following equation:

$$R(S \mid D) = \frac{\sum_{x \in C} \sum_{b \in W} Freq(H, b)}{D} \tag{6}$$

where $R(S \mid D)$ is the representativeness of Sample Set S and D is the complete dataset. H is the histogram of D_p. $Freq(H, b)$ is the number of points in (or the frequency of) any bucket b in histogram H. The window size W that is a set of buckets b. A selected region represents all regions in W. W can be user-defined. C is the set of regions or centroids selected.

4.2.1 Preconditions

Before any of the following workflows are executed, the user selects a subset of the global variables (such as temperature, %tree cover) for which representativeness is measured. Also, the user is allowed to filter out a set of regions from the dataset. For example, the user may choose to limit the dataset only to regions which are in the tropics (i.e., having a temperature between 15^oC and 31^oC). Once the data is filtered, we get a final list of regions and matrix containing the value of each user selected variable for each of these regions.

4.2.2 Workflow 1 - Calculate representativeness of given sample set

The steps to calculate representativeness of a sample set of sites (equation 6) that are provided by the user are:

Step 1

Select the a set of variables based on which representativeness is measured. Calculate the eigenvalues and eigenvectors for all the unfiltered regions across the world. If the unfiltered regions do not contain the sites that are part of the Sample Set, then find the first principal component and project all the regions (including the sites) onto this dimension.

Step 2

Calculate the distance between each region and a region in the Sample Set. The Final Distance is the minimum distance found. It shows which location in the Sample Set is closest to the current location being considered and how close it is.

Step 3

Display all the locations on Google Maps with a color chosen according to the final distance calculated. All the distance are a value between 0 & 1. A color pallete consists of a set of shades (e.g. from Green to Red where Green is considered as a place completely represented by one of the sample sites while Red is completely unrepresented). The scale between 0 & 1 is divided based on the number of shades which are in the color pallete. The color associated with the bucket in which the final distance of a place falls, is applied on Google Maps, on top of the location of that place. The following diagram (figure 4.2.2) is an example of how the color scheme looks.

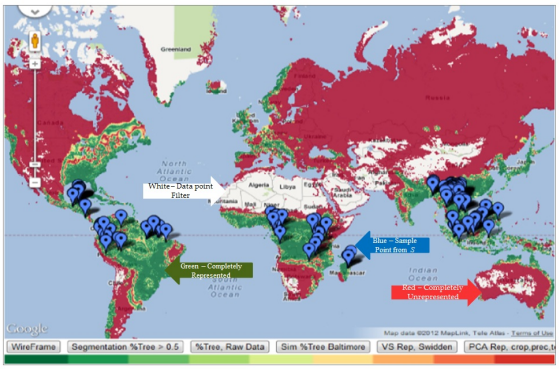

Figure 5: *Representativeness Coloring Scheme*

4.2.3 Workflow 2 - Creating a set of *Ideal* sites

An *ideal* set of sites is a set of sites that maximizes representativeness but at the same time has a minimum number of sites required to reach that measure of representativeness. The steps to find a set of *ideal* set of sites that are representative of all other regions in the world are:

Step 1

As seen in the previous workflow, calculate the eigenvalues and eigenvectors for all the unfiltered regions across the world.

Step 2

Once the regions are projected onto the first principal component, create the histogram of that dimension. The histogram requires a number of buckets into which the dimension is divided. The

number of buckets is defined after testing. Once the histogram is constructed, a window of size *W* is defined such that

$$1 \leq \text{W} \leq \textit{number of buckets}$$

The window is a set of buckets (containing regions) such that any region in the bucket represents all the regions in the buckets which are present in *W*. *W* is either user-defined or an arbitrarily fixed size. The algorithm performs the following steps *N* times where *N* is the number representative sites to be found.

Algorithm 1: Finds M points to maximize Representativeness of Points

Input: A finite set $FB_P = \{fb_1, fb_2, \ldots, fb_m\}$ of Frequency of buckets in Histogram H, window size W
Output: A finite set C containing a set of N representative Points
rand() is picks a random point from a bucket b_k
WUsed is bit array of size $\lfloor \frac{B_P}{W} \rfloor$
B_P contains points each bucket. $B_P = \{b_1, b_2, \ldots, b_m\}$
MW is the window with Maximum Frequency.

```
for i ← 1 to ⌊B_P / W⌋ do
    WUsed_i ← false
for i ← 1 to N do
    max ← −1
    j ← 1
    MW ← −1
    maxBin ← −1
    while k ≤ n do
        if (fb_k ≥ max) && (WUsed_j = false) then
            max ← fb_k
            maxBin ← b_k
            MW ← j
        if mod(k, W) = 0 && (k/W) ≥ 0 then
            j ← j + 1
        k ← k + 1
    WUsed_MW ← true
    P ← rand(b_maxBin)
    C_i ← P
return C
```

Algorithm 1 shows how an *ideal* set of sites is selected. Consider a histogram *H*, with window size *W*. Let B_P contains points each bucket. MW is the window with the maximum frequency. $\lfloor \frac{B_P}{W} \rfloor$ is the total number of windows in histogram *H*. We maximize representativeness by selecting a single point from a window. Once a window is utilized, it not used again. This is because a point from a given window of buckets represents all the points in the window completely. *WUsed* is an array of bits that shows which of the windows have used. *WUsed* is initialized to $false$. Then, the algorithm iterates through each bucket in the histogram and finds the bucket which has the maximum frequency (mode) and that is part of a window that has not been used before (i.e., where $WUsed_i$ is $false$) . A bucket b_k is part of a window $\lfloor (k/W) \rfloor$. If bucket b_k is the bucket with the maximum freqeuncy, then make $WUsed_{\lfloor (k/W) \rfloor}$ as $true$. Then, select any point from b_k and add it to the final list of centroids *C*. The above set of steps performed to find a centroid, is repeated *N* times to get a set *C* containing *N* centroids.

The algorithm has to take care of a specific condition i.e. the case when the number of centroids required is higher as compared to the number of bins that are greater than 0. In such a case, the number of centroids returned is limited to number of non-zero bins that can be found.

Step 3
This step is the same as the previous workflow, where the regions selected are displayed on a World Map with the representativeness of the entire set.

5 EXPERIMENTAL RESULTS & ANALYSIS

Land Change scientists also compare their sample sites against random sampling, as a standard practice. Each experiment conducts 3 sets of tests. The first test is to calculate the representativeness of the sample sites given by the user. The second test uses the *histogram* method discussed in the previous chapter to generate a new set of sites (the number of sites is equal to that defined by the user) where a study can be conducted. It also calculates the representativeness for the same. The third test generates representativeness for a random set of sites for the same number of sites as in the previous two cases. Random sampling is performed 1000 times to eliminate any bias created from a limited set of random sampling tests. We define the ideal sample set as the set of samples generated by the histogram method 1.

5.1 Measuring Representativeness

The Van Vliet study has been used for the purpose of conducting experiments van Vliet et al. [2012]. This meta-study conducts a global assessment of swidden cultivation i.e. slash and burn. It is an agricultural technique where forested areas are burned to create fields for agriculture. There are a total of 157 sites that are part of the study. These sites are used as centroids in our experiment to measure its representativeness. The window size for all experiments below is 1 and number of bins in the histogram is 157.

Figure 6 shows the location of sites that are part of the Van Vliet Study.

Figure 6: *Sites in Van Vliet Study*

5.1.1 Measuring with a Filter

The following parameters are applied in the experiment:
Filter is applied so that regions of the world the author does not claim to represent are not considered in the analysis. In the example, the parameter *potential vegetation* is used to filter the data set. Potential vegetation has a range of values form 0 to 12 Ramankutty [1999]. The values considered in the experiment are from 1 to 2. These values are used filter out all regions except the tropical regions (and some forested areas) across the world.

A total of 3 attributes or variables are used in the experiment viz. *potential vegetation*, *market access* and *temperature*. The selection of the variables is based on the study.

Zoom Level decides the total number of points which are considered in the dataset. For example, at zoom level 4, the total number of points or regions is 160,000. The experiments are conducted at zoom level 6. The total number of points at this level is about 1.4 million. When the filter is applied, the total number of *Global land units GLU's* is reduced to about 250,000.

Color Scale applied has a total of 10 colors from red to green. It depicts the different levels of representation in the *heat map*.

The figures 7, 8, and 9 depict the *heat map* generated for the Van Vliet study. The color scale is provided at the bottom of each image. The scale goes from red to green, where green depicts “complete represention" and red depicts “complete non-representation". The parts of the world that are filtered out are shown in dark blue. This contains all the water bodies and the other regions of the world that do not fit the filter criteria.

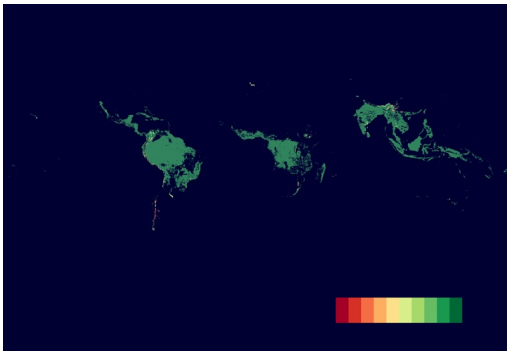

Figure 7: *Representativeness of Filtered World Regions Using Given Samples*

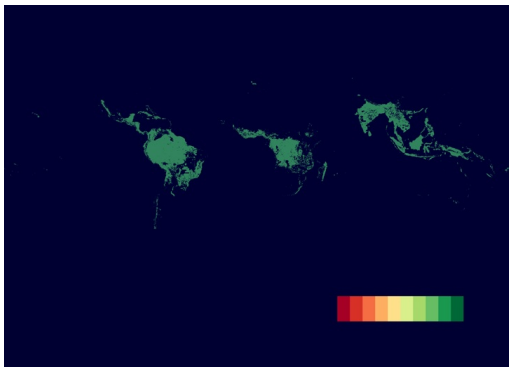

Figure 8: *Representativeness of Filtered World Regions Using Ideal Samples*

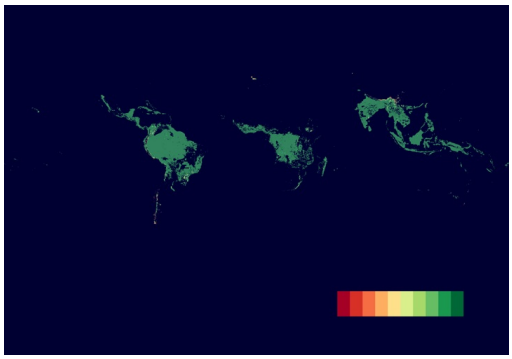

Figure 9: *Representativeness of Filtered World Regions Using Random Samples*

Figure 7 shows some regions that are not represented by the locations in the study. The maps generated using given, random and ideal samples are almost the same. This is because of the definition of representativeness. To maximize representativeness, the sampled regions have to “cover" as much of the filtered regions as possible. To represent a set of regions that have similar conditions, only a

single point or location is required. Based on the filter used, the potential vegetation, temperature and market access parameters are of similar areas as can be seen on the map. This makes regions cluster with high frequency in specific regions in the PCA space, thus making the number of locations required to represent them lower than 157. This is seen in the histogram in figure 10 where the high frequency bins require only a single location in the bin to represent the entire bin. Thus random sampling and the histogram method to calculate new locations have representativeness values that are close to *1.0* or 100%.

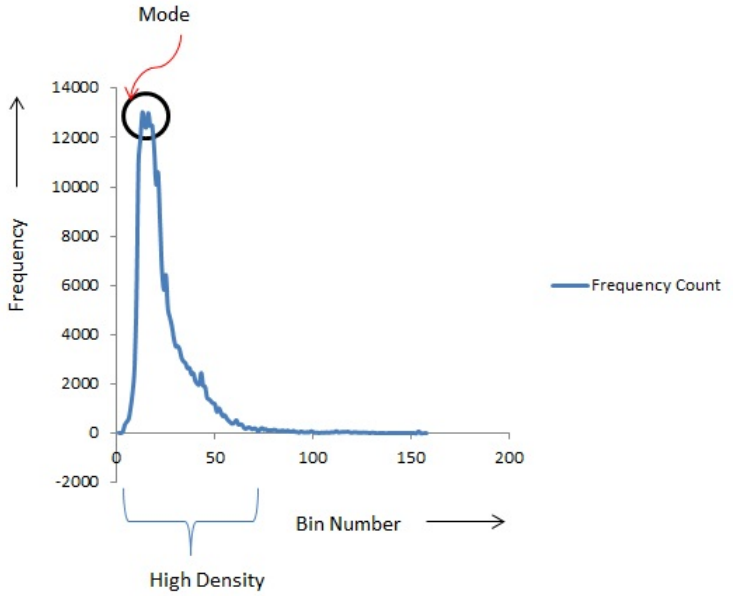

Figure 10: *Histogram Of First Principal Component Values (For Filtered Data) of 1.4 Million Regions*

Figure 12 shows the comparison of representativeness of the given sample and the ideal samples generated against random sampling conducted. Representativeness for the given sample set at 48^{th} percentile while the ideal sample set is at 100 percentile. It means that the given sample set is better than random sampling only 48% of the time as compared to ideal sample set which is better 100% of the time. This makes the ideal sample set a better selection as compared to random or the given sample set. Table 1 shows that representativeness of the given sample set is better the representativeness of the average random sample set.

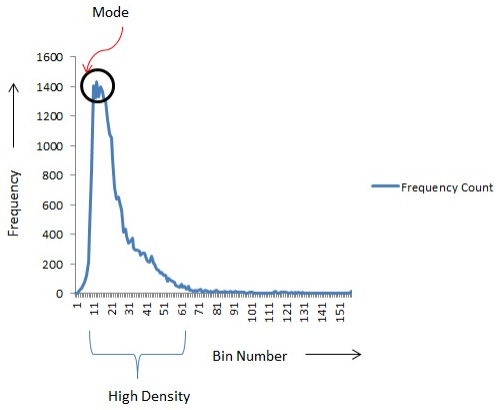

Figure 11: *Histogram Of First Principal Component Values (For Filtered Data) of 160000 Regions*

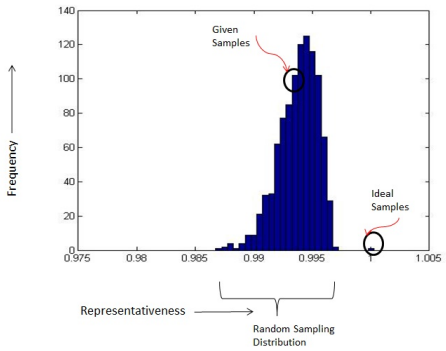

Figure 12: *Histogram Of Representativeness for Random Sampling and Where Other Methods Lie*

Method	Representativeness
Given Sample	0.994
Ideal Sample	1.0
Avg. Random Sampling	0.9937

Table 1: *Representativeness Of Samples for Filtered Data*

Figure 11 shows that the same histogram trend is maintained when we perform PCA on a downsized dataset. Hence the calculations made on a ISea3H level 10 are applied to the ISea3H level 12 hexagons. Thus representativeness, based on a downsized number of regions, does not change as nature of distribution remains the same. The total number of regions for the downsized set after filtering is 27883.

5.1.2 Measuring without a Filter

This section describes an experiment to compare the various samples without a filter being applied. The data set is the complete 1.3 million regions across the globe. The world maps generated show a clearer differentiation in the given sample representativeness against that of an ideal and a random sample.

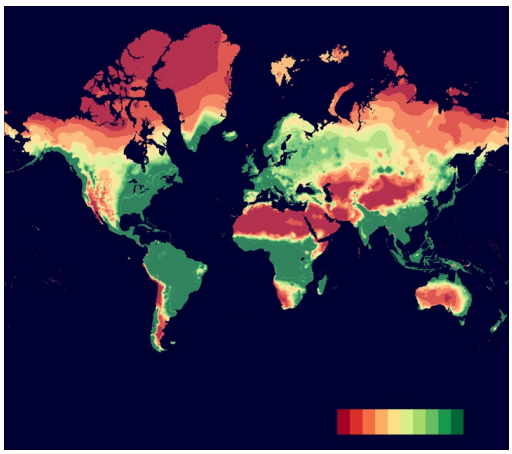

Figure 13: *Representativeness of Unfiltered World Regions Using Given Samples*

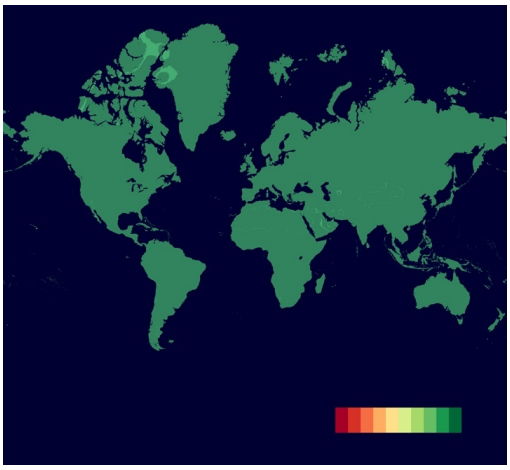

Figure 14: *Representativeness of Unfiltered World Regions Using Ideal Samples*

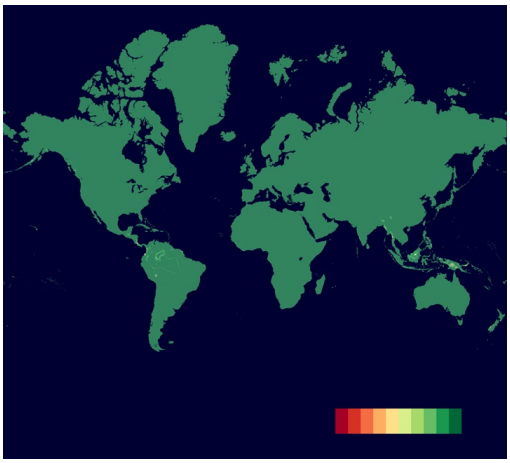

Figure 15: *Representativeness of Unfiltered World Regions Using Random Samples*

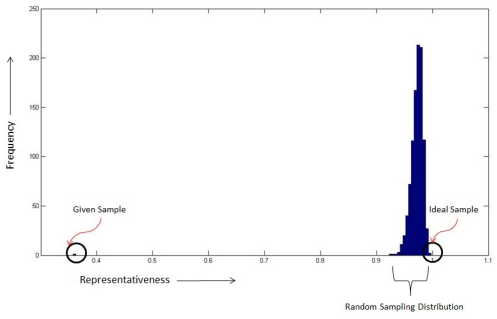

Figure 16: *Histogram Of Representativeness for Random Sampling and Where Other Methods Lie*

Figure 16 gives an idea of the distribution of representativeness values for random sampling and shows where the given sample and ideal samples are placed in the distribution. Table 2 shows that ideal sampling has the highest representativeness.

Method	Representativeness
Given Sample	0.362
Ideal Sample	0.995
Avg. Random Sampling	0.972

Table 2: *Representativeness Of Samples for Unfiltered Data*

The variables selected affect the nature of the distribution. A uniform distribution of points across variables will reduce the amount variation covered by a single principal component. For example, if the variables distributed the data points in the form of a circle (in a 2 dimensional graph), then there would not be a one principal component that would cover a large portion of the variation limiting the use of PCA. The variables we have used in our experiments have a large variation (e.g. temperature) where as potential vegetation is a categorical value limiting variation.

5.2 Ideal vs Random Sampling

In the previous section, the results show that ideal sampling has a representativeness close to 1.0, the same as random sampling. This result, as described previously, is a function of the number of sites (or centroids) that need to be generated. Since the number of sites (157) is large, the representativeness is close to *1.0*. We can analyze the effectiveness of ideal sampling by measuring the representativeness of ideal samples against random sampling for a reduced of sample size. We consider the unfiltered data points as in previous example. The following graph shows the trend for representativeness for the two methods, starting from the selection of 1 sample site to a total of 157.

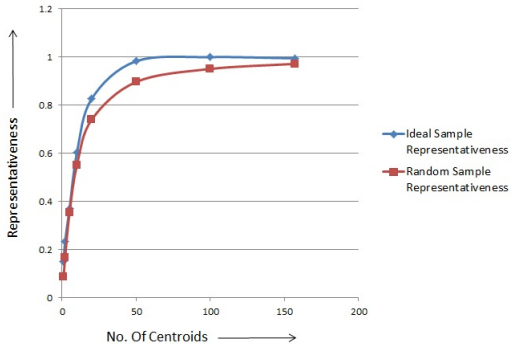

Figure 17: *Representativeness Trend - Increasing Number Of Centroids*

The trend shows that ideal sample representativeness tends to 1 to with fewer centroids as compared to random sampling, making it better to select sites as compared to random sampling. Table 3 shows the number of centroids required by each method to reach a value for representativeness close to 1.

Method	Centroids	R
Ideal	60	0.99
Avg. Of Random	100 - 130	0.95 - 0.967

Table 3: *Number Of Centroids & Representativeness*

5.3 Measuring the Effect of Histogram Size

The ideal samples are generated by constructing a histogram of the first principal component. The algorithm 1 is affected by 2 parameters, i.e., the number of bins created in the histogram and the window size. The window size is the number bins in which any data point can be considered to represent the rest of the data points with the bins in the window completely. Currently, the window size and number of bins are chosen after testing. Figure 18 shows the change in representativeness for various window sizes and varying number of bins.

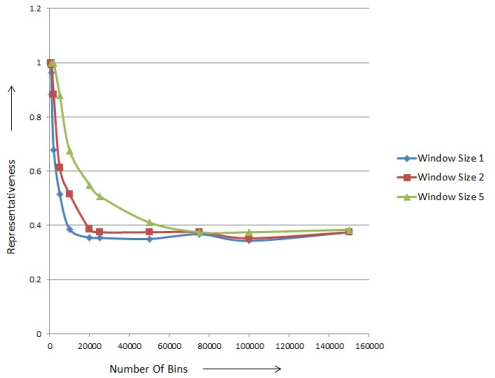

Figure 18: *Representativeness Trend - Increasing Number Of Bins*

The test has been conducted for window sizes 1, 2 & 5. The minimum number of bins taken is equal to the number of centroids required which is maintained at 157 throughout all the tests. As the window size is 1, there are a total of 157 bins from which a data point can be selected (refer to the algorithm 1). As the number of bins is increased, the representativeness decreases, until it reaches a set of bins after which the representativeness becomes stable. The second observation is that the rate of decrease in representativeness is lower when the window size is increased. Thus a higher window size can reduce rate of decrease in representativeness. We explain the reasons for the decrease in the next section.

5.3.1 Effect of the Number of Bins

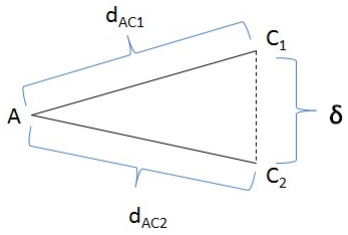

Figure 19: *Example for Bin Distance*

Consider figure 19. There are 3 points $A, C_1,$ *and* C_2, where $C_1 \& C_2$ are 2 centroids. $d_{AC1} \& d_{AC2}$ are the distances between the points $A \& C_1$ and $A \& C_2$ in the first principal component respectively. δ is the distance between $C_1 \& C_2$. From the diagram, we see if $\delta \to 0$, then $(d_{AC1} - d_{AC2}) \to 0$.

$$\textit{IntervalSize} = \frac{p_{max} - p_{min}}{NoOfBins} \tag{7}$$

where p_{max} and p_{min} are the maxmimum and minimum values in the first principal component respectively.

Thus, when the number of bins increases, IntervalSize decreases. Consider the case when window size is 1. Consider figure 10 where there are certain parts of the histogram that have high density. As the IntervalSize decreases, centroids are chosen from bins adjacent to the mode where $\delta \to 0$. The final representativeness is based on the color scale that is used. All the data points that fall in the first bin of the scale are considered completely represented by the sample set. As more centroids are selected, the difference in the distance either remains the same or get smaller. The redundant centroids generated are unable to *cover* the rest of the points in the distribution to minimize the distance and maximize coverage. Thus,

$$NoOfBins \; \alpha \; Colorscale \tag{8}$$

Representativeness stabilizes after a certain number of bins because the histogram is divided into small parts such that coverage is only for the high density part of the histogram from where all the centroids are selected. In such a case, increasing the number of bins no longer affects the representativeness.

5.3.2 Effect of the Window Size

The window size creates a minimum distance between any 2 centroids that are selected as only a single point can be selected from within a bin in a certain window at a time. Hence figure 18 shows a lower rate of decrease in representativeness as the number of bins increase. For a optimal representativeness, the ideal sample set selected should be function of the number of bins and size of the window used. Hence equation 8 can be modified to,

$$\frac{NoOfBins}{WindowSize} \; \alpha \; Colorscale \tag{9}$$

6 CONCLUSION & FUTURE WORK

We have provided a definition and an algorithm for calculating the representativeness of a set of sample sites. When the number of dimensions increases, clustering methods become computationnally inefficient. This is specifically when the representativeness needs to be calculated in near realtime. Hence dimension reduction techniques are used. We have used principal component analysis to perform n-dimension reduction into a single dimension based on the variance of attributes. This helps us project points onto the first principal component with maximum difference or spacing between points. To nullify the effect of the magnitude of the values in each attribute, normalization of each column is performed. This makes all the values fall in the scale from 0 to 1. The distances from the given sample set to other points are calculated using their first principal component projected values. We see that the results of the heat map generated to show representativeness across the globe is as expected. To maximize representativeness, we have provided a method that is based on the creating a histogram of the PCA values and selecting modes. The method is able to maximize representativeness as seen in the experiments conducted. The samples drawn from land change literature and the ideal samples are compared against random sampling. We show that if a filter is applied, the given sample set has the lowest representativeness as compared to the other two methods (figure 12 and figure 16), even though representativeness is greater than 0.9. It also shows that ideal sampling is better than random sampling. We are also able to see, as shown in figure 17, that ideal sampling reaches the same measure of representativeness as compared to random sampling with fewer samples. Hence it is better than random sampling at performing site selection. We also analyze the properties of the ideal sampling method for each of the parameters affecting the method, mainly the number of bins in the histogram and the window size. We show (figure 18) that when the number of bins is increased the representativeness decreases until it reaches a stable level.

In the future, a set of improvements can be performed. These are:

1. Creating a function correlation between the number of bins in the histogram to the scale applied for representativeness. The current method applies an arbitrary number of bins in the histogram. This helps to maximize representativeness under all conditions.
2. The scale applied is a linear scale from 0 to 1 that is divided equally. There is no correlation between the scale and the actual distance calculated between the centroids and the other points. If the initial projected values are very small, then the distances calculated are also small. Thus a region that is not related to any of the centroids can be shown as being represented. One solution to this problem is to normalize the projected values again, so that in case the values are very small they are scaled up accordingly to a value between 0 and 1. The points that have a value greater than 1 can be considered as outliers.
3. The problem of the distance (between a sample and another region in PCA space) being very small may still exist. The representativeness scale can be changed to account for regions in the principal component where the density of points is concentrated. Thus, density

estimation and PCA outlier detection can be performed to create a tighter lower and upper bound of projected data points.

4. The visualization of the map can be improved by implementing isolines or contour lines. Contour lines or isolines are lines across which the function has the same output value. Thus isolines can be implemented for all the areas that have the same projected value on the first principal component.

References

H. Abdi and L. J. Williams. Principal component analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(4):433–459, 2010. ISSN 1939-0068. doi: 10.1002/wics.101. URL `http://dx.doi.org/10.1002/wics.101`.

D. Brugger, M. Bogdan, and W. Rosenstiel. Automatic cluster detection in kohonen's som. *Neural Networks, IEEE Transactions on*, 19(3):442 –459, march 2008. ISSN 1045-9227. doi: 10.1109/TNN.2007.909556.

S. Bubeck, U. V. Luxburg, and C. Elkan. Nearest neighbor clustering: A baseline method for consistent clustering with arbitrary objective functions, 2009.

M. A. Carreira-Perpinan. A review of dimension reduction techniques. Technical report, Dept. of Computer Science, University of Sheffield, January 27 1997a. Technical Report CS-96-09.

M. A. Carreira-Perpinan. A review of dimension reduction techniques. Technical Report CS-96-09, Department of Computer Science, University of Sheffield, 1997b.

C. Ding and X. He. K-means clustering via principal component analysis. In *Proceedings of the twenty-first international conference on Machine learning*, ICML '04, pages 29–, New York, NY, USA, 2004. ACM. ISBN 1-58113-838-5. doi: 10.1145/1015330.1015408. URL `http://doi.acm.org/10.1145/1015330.1015408`.

E. Ellis. Cdi-type ii: Globe: Evolving new global workflows for land change science, 2011.

R. T. M. S. M. D. J. E. Forrest M. Hoffman, William W. Hargrove and R. J. Oglesby. Multivariate spatio-temporal clustering (mstc) as a data mining tool for environmental applications. *In Miquel Sanchez-Marre, Javier BAejar, Joaquim Comas, Andrea E. Rizzoli, and Giorgio Guariso, editors, Proceedings of the iEMSs Fourth Biennial Meeting: International Congress on Environmental Modelling and Software (iEMSs 2008)*, Barcelona, Catalonia, Spain, July 2008.

A. K. Jain, M. N. Murty, and P. J. Flynn. Data clustering: a review. *ACM Comput. Surv.*, 31(3): 264–323, Sept. 1999. ISSN 0360-0300. doi: 10.1145/331499.331504. URL `http://doi.acm.org/10.1145/331499.331504`.

C. D. Manning, P. Raghavan, and H. Schtze. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008. ISBN 0521865719, 9780521865715.

I. P. on Climate Change. Summary for policymakers - land use, land-use change, and forestry. Technical report, Intergovernmental Panel on Climate Change, 2000.

B. C. Pijanowski, D. G. Brown, B. A. Shellito, and G. A. Manik. Using neural networks and gis to forecast land use changes: a land transformation model. *Computers, Environment and Urban Systems*, 26(6):553 – 575, 2002. ISSN 0198-9715. doi: 10.1016/S0198-9715(01)00015-1. URL `http://www.sciencedirect.com/science/article/pii/S0198971501000151`.

J. F. Ramankutty, N. Estimating historical changes in global land cover: croplands from 1700 to 1992. *Global Biogeochemical Cycles*, 13(4):997–1027, 1999.

L. T. Tran, C. G. Knight, R. V. ONeill, E. R. Smith, and M. OConnell. Self-organizing maps for integrated environmental assessment of the mid-atlantic region. *Environmental Management*, 31: 822–835, 2003. ISSN 0364-152X. URL `http://dx.doi.org/10.1007/s00267-003-2917-6`. 10.1007/s00267-003-2917-6.

N. van Vliet, O. Mertz, A. Heinimann, T. Langanke, U. Pascual, B. Schmook, C. Adams, D. Schmidt-Vogt, P. Messerli, S. Leisz, J.-C. Castella, L. JÃ¸rgensen, T. Birch-Thomsen, C. Hett, T. Bech-Bruun, A. Ickowitz, K. C. Vu, K. Yasuyuki, J. Fox, C. Padoch, W. Dressler, and A. D. Ziegler. Trends, drivers and impacts of changes in swidden cultivation in tropical forest-agriculture frontiers: A global assessment. *Global Environmental Change*, 22(2):418 – 429, 2012. ISSN 0959-3780. doi: 10.1016/j.gloenvcha.2011.10.009. URL `http://www.sciencedirect.com/science/article/pii/S0959378011001622`. Adding Insult to Injury: Climate Change, Social Stratification, and the Inequities of Intervention.

A. Vattani. k-means requires exponentially many iterations even in the plane. In *Proceedings of the 25th annual symposium on Computational geometry*, SCG '09, pages 324–332, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-501-7. doi: 10.1145/1542362.1542419. URL `http://doi.acm.org/10.1145/1542362.1542419`.